UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/748,410	12/29/2003	Fay Chong JR.	188164/US	7553

66083   7590   09/16/2008
SUN MICROSYSTEMS, INC. c/o DORSEY & WHITNEY, LLP
370 SEVENTEENTH ST.
SUITE 4700
DENVER, CO 80202

EXAMINER
LE, MIRANDA

ART UNIT	PAPER NUMBER
2169	

MAIL DATE	DELIVERY MODE
09/16/2008	PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE <u>3</u> MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.
- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on <u>04 June 2008</u>.

2a)☒ This action is **FINAL**.   2b)☐ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) <u>1,2,4-16,18-30 and 32-46</u> is/are pending in the application.

    4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) <u>1,2,4-16,18-30 and 32-46</u> is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

    Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

    Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11)☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

    a)☐ All   b)☐ Some * c)☐ None of:

      1.☐ Certified copies of the priority documents have been received.

      2.☐ Certified copies of the priority documents have been received in Application No. _____.

      3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

    * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) ☒ Notice of References Cited (PTO-892)

2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)

3) ☒ Information Disclosure Statement(s) (PTO/SB/08) Paper No(s)/Mail Date <u>06/04/08, 03/14/08</u>.

4) ☐ Interview Summary (PTO-413) Paper No(s)/Mail Date. _____ .

5) ☐ Notice of Informal Patent Application

6) ☐ Other: _____.

DETAILED ACTION

This communication is responsive to Amendment, filed 06/04/08.

Claims 1-2, 4-16, 18-30, 32-46 are pending in this application. Claims 1,
15, 29, 43 are independent claims. This action is made non-Final.

Information Disclosure Statement

Applicants' Information Disclosure Statements, filed 06/04/2008,
03/14/2008, have been received, entered into the record, and considered. See
attached form PTO-1449.

Specification

The objection to the specification of the invention has been withdrawn in
view of the amendment.

Claim Rejections - 35 USC § 101

The rejection of claims 15-28 under 35 U.S.C. §101 has been withdrawn
in view of the amendment. It is acknowledged that Applicant has amended the
specification by deleting on page 2, paragraph 0042, the reference to "a
propagated signal", "carrier wave". Thus, Applicant has provided evidence that "a
machine readable medium" intended to be covered only as "ROM, RAM, disk
storage media". As such, claims 15-28 recite an only statutory subject matter.

Claim Rejections - 35 USC § 112

The following is a quotation of the second paragraph of 35 U.S.C. 112:

The specification shall conclude with one or more claims particularly pointing out and distinctly claiming the subject matter which the applicant regards as his invention.

Claims 29-46 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention.

Claims 29-46 recite the limitation of "receiving a command to create a backup". The specification does not provide a standard for ascertaining the requisite degree, and one of ordinary skill in the art would not be reasonably appraised of the scope of the invention. Therefore the limitation is indefinite as it fails to point out what is being described since the specification does not clearly redefine the term.

Appropriate correction is required.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

This application currently names joint inventors. In considering patentability of the claims under 35 U.S.C. 103(a), the examiner presumes that the subject matter of the various claims was commonly owned at the time any inventions covered therein

were made absent any evidence to the contrary. Applicant is advised of the

obligation under 37 CFR 1.56 to point out the inventor and invention dates of each

claim that was not commonly owned at the time a later invention was made in order

for the examiner to consider the applicability of 35 U.S.C. 103(c) and potential 35

U.S.C. 102(e), (f) or (g) prior art under 35 U.S.C. 103(a).

Claims 1, 2, 6-8, 11, 12, 15, 16, 20-22, 25, 26, 29, 30, 34-36, 39, 40, 43-

46 are rejected under 35 U.S.C. 103(a) as being unpatentable over Paulson

et al. (US Patent No. 6,112,319), in view of Flynn et al. (US Patent No.

6,453,392), and further in view of Tzelnic et al. (US Patent No. 6,366,987).

As per claim 1, Paulson teaches a method for preserving data in a data

storage system, the method comprising:

receiving a command to write data to the data storage system *(i.e. a*

second write request ("2w") to write the value 2, a second read request ("2f"), col.

4, lines 45-50);

executing *(i.e. a first write request ("1w") to write the value 1, a first read*

("1r") request, col. 4, lines 45-50), for a first data, a fist input/output (I/O) process

directed to a first storage volume, wherein the first storage volume is not mirrored

and the first I/O process *(i.e. a first write request ("1w") to write the value 1, a first*

read ("1r") request, col. 4, lines 45-50) begins at a first time which is prior to

receiving the command *(i.e. a second write request ("2w") to write the value 2, a*

second read request ("2f"), col. 4, lines 45-50);

creating a data structure, in response to the command, for at least a

second image, the second image storing changes to the first storage volume

occurring after receipt of the command *(i.e. it creates a new data cell, stores the data from the write request in the new cell, and appends the new data cell to the end of the data list for this address cell, col. 7, line 56 to col. 8, line 23)*;

writing a second data *(i.e. When the second write request is received, another data cell would be created that held the value of 2, col. 7, line 56 to col. 8, line 23)* directed to the second image as part of a second I/O process *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50)* which begins after receipt of the command *(i.e. it creates a new data cell, stores the data from the write request in the new cell, and appends the new data cell to the end of the data list for this address cell, col. 7, line 56 to col. 8, line 23)*; and

modifying the data structure *(i.e. modifies the data structure to reflect the responses, col. 8, lines 55-62)* to indicate that the second data is stored in the second image and storing the second data in the second image *(i.e. If it is for a write request, the routine continues to step 329 where it retrieves the data list for the address cell and removes all of the data cells from the list before the data cell corresponding to the write, col. 8, line 64 to col. 9, line 9)*.

Paulson does not specifically teach:

receiving a command to create <u>a backup of data stored in the</u> data storage system;

a second storage volume;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command; and

Flynn teaches:

receiving a command to access data to the data storage system *(i.e. In a virtual machine environment, a virtual machine ID associated with a virtual machine is sent to a storage controller along with a <u>request to access a storage device</u>. The virtual machine ID is used by the storage controller along with a path group ID to determine the scope of the access request. Thus, virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally, Summary)*;

a second storage volume *(i.e. Accordingly, a storage controller is operative to deny a storage device access request from a first virtual machine of a host processor responsive to a compare by the storage controller of a first virtual machine ID of the first virtual machine and a second virtual machine ID of a second virtual machine of the host processor. More specifically, the storage controller may be operative to receive a request for accessing a storage device; operative to compare a stored path group ID and a requesting path group ID associated with the request; operative to compare a stored virtual machine ID and a requesting virtual machine ID associated with the request; and operative to grant or deny the request responsive to the compares, col. 10, lines 24-37)*;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command *(i.e. the guest operating system performs device operations (read, write, etc.) on storage device 128, col. 8, lines 24-36)*.

It would have been obvious to one of ordinary skill of the art having the teaching of Paulson and Flynn at the time the invention was made to modify the system of Paulson to include the limitations as taught by Flynn. One of ordinary skill in the art would be motivated to make this combination in order to determine the scope of the access request in view of Flynn (Summary), as doing so would give the added benefit of virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally as taught by Flynn (Summary).

Paulson and Flynn do not explicitly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage system;

Tzelnic teaches this limitation *(i.e. <u>In response to the back-up request</u>, the method includes selecting a physical storage unit containing the logical data structure, and sending a backup command to the primary data storage subsystem for backing up the physical storage unit by copying a current version of the physical storage unit from primary storage in the primary data storage subsystem to the secondary data storage subsystem <u>to create a backup version</u> of the physical storage unit in secondary storage in the secondary data storage subsystem, col. 1, line 58 to col. 2, line 17).*

It would have been obvious to one of ordinary skill of the art having the teaching of Paulson, Flynn, and Tzelnic at the time the invention was made to modify the request for accessing a storage system of Paulson, Flynn to include a request for creating a backup of data in a storage system as taught by Tzelnic.

One of ordinary skill in the art would be motivated to make this combination in order to restore the logical data structure in view of Tzelnic (col. 1, line 58 to col. 2, line 17), as doing so would give the added benefit of restoring the data rapidly from backup after the storage system failure as taught by Tzelnic (col. 1, lines 33-55).

As per claim 15, Paulson teaches an article of manufacture comprising:

a machine-readable medium having executable code to cause a machine to perform a method for preserving data in a data storage system *(See Fig. 1)*, the method comprising:

receiving a command to write data to the data storage system *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50)*;

executing *(i.e. a first write request ("1w") to write the value 1, a first read ("1r") request, col. 4, lines 45-50)*;

executing *(i.e. a first write request ("1w") to write the value 1, a first read ("1r") request, col. 4, lines 45-50)*, for a first data, a fist input/output (I/O) process *(i.e. a first write request ("1w") to write the value 1, a first read ("1r") request, col. 4, lines 45-50)* directed to a first storage volume, wherein the first storage volume is not mirrored and the first I/O process begins at a first time which is prior to receiving the command *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50)*;

creating a data structure, in response to the command, for at least a

second image, the second image storing changes to the first storage volume

occurring after receipt of the command *(i.e. it creates a new data cell, stores the*

data from the write request in the new cell, and appends the new data cell to the

end of the data list for this address cell, col. 7, line 56 to col. 8, line 23);

writing a second data *(i.e. When the second write request is received,*

another data cell would be created that held the value of 2, col. 7, line 56 to col.

8, line 23) directed to the second image as part of a second I/O process *(i.e. a*

second write request ("2w") to write the value 2, a second read request ("2f"), col.

4, lines 45-50) which begins after receipt of the command *(i.e. it creates a new*

data cell, stores the data from the write request in the new cell, and appends the

new data cell to the end of the data list for this address cell, col. 7, line 56 to col.

8, line 23); and

modifying the data structure *(i.e. modifies the data structure to reflect the*

responses, col. 8, lines 55-62) to indicate that the second data is stored in the

second image and storing the second data in the second image *(i.e. If it is for a*

write request, the routine continues to step 329 where it retrieves the data list for

the address cell and removes all of the data cells from the list before the data cell

corresponding to the write, col. 8, line 64 to col. 9, line 9).

Paulson does not fairly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage

system;

a second storage volume;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command; and

Flynn teaches:

receiving a command to access data to the data storage system *(i.e. In a virtual machine environment, a virtual machine ID associated with a virtual machine is sent to a storage controller along with a <u>request to access a storage device</u>. The virtual machine ID is used by the storage controller along with a path group ID to determine the scope of the access request. Thus, virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally, Summary)*;

a second storage volume *(i.e. Accordingly, a storage controller is operative to deny a storage device access request from a first virtual machine of a host processor responsive to a compare by the storage controller of a first virtual machine ID of the first virtual machine and a second virtual machine ID of a second virtual machine of the host processor. More specifically, the storage controller may be operative to receive a request for accessing a storage device; operative to compare a stored path group ID and a requesting path group ID associated with the request; operative to compare a stored virtual machine ID and a requesting virtual machine ID associated with the request; and operative to grant or deny the request responsive to the compares, col. 10, lines 24-37)*;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command *(i.e. the guest*

operating system performs device operations (read, write, etc.) on storage device

128, col. 8, lines 24-36).

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson and Flynn at the time the invention was made to modify the

system of Paulson to include the limitations as taught by Flynn. One of ordinary

skill in the art would be motivated to make this combination in order to determine

the scope of the access request in view of Flynn (Summary), as doing so would

give the added benefit of virtual machines of a single host processor may share

the storage device while both preserving data integrity and performing optimally

as taught by Flynn (Summary).

Paulson and Flynn do not explicitly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage

system;

Tzelnic teaches this limitation *(i.e. <u>In response to the back-up request</u>, the*

method includes selecting a physical storage unit containing the logical data

structure, and sending a backup command to the primary data storage

subsystem for backing up the physical storage unit by copying a current version

of the physical storage unit from primary storage in the primary data storage

subsystem to the secondary data storage subsystem <u>to create a backup version</u>

of the physical storage unit in secondary storage in the secondary data storage

subsystem, col. 1, line 58 to col. 2, line 17).

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson, Flynn, and Tzelnic at the time the invention was made to

modify the request for accessing a storage system of Paulson, Flynn to include a

request for creating a backup of data in a storage system as taught by Tzelnic.

One of ordinary skill in the art would be motivated to make this combination in

order to restore the logical data structure in view of Tzelnic (col. 1, line 58 to col.

2, line 17), as doing so would give the added benefit of restoring the data rapidly

from backup after the storage system failure as taught by Tzelnic (col. 1, lines

33-55).

       As per claim 29, Paulson teaches an apparatus for preserving data in a

data storage system, comprising:

       a processing system configured to process computer-executable code

(See Fig. 1), the computer-executable code causing the processing system to

perform the operations of:

       receiving a command to write data to the data storage system *(i.e. a*

second write request ("2w") to write the value 2, a second read request ("2f"), col.

4, lines 45-50);

       executing *(i.e. a first write request ("1w") to write the value 1, a first read*

("1r") request, col. 4, lines 45-50);

       means for executing *(i.e. a first write request ("1w") to write the value 1, a*

first read ("1r") request, col. 4, lines 45-50), for a first data, a fist input/output

(I/O) process directed to a first storage volume, wherein the first storage volume

is not mirrored and the first I/O process *(i.e. a first write request ("1w") to write*

the value 1, a first read ("1r") request, col. 4, lines 45-50) begins at a first time

which is prior to receiving the command *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50)*;

creating a data structure, in response to the command, for at least a second image, the second image storing changes to the first storage volume occurring after receipt of the command *(i.e. it creates a new data cell, stores the data from the write request in the new cell, and appends the new data cell to the end of the data list for this address cell, col. 7, line 56 to col. 8, line 23)*;

writing a second data *(i.e. When the second write request is received, another data cell would be created that held the value of 2, col. 7, line 56 to col. 8, line 23)* directed to the second image as part of a second I/O process *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50)* which begins after receipt of the command *(i.e. it creates a new data cell, stores the data from the write request in the new cell, and appends the new data cell to the end of the data list for this address cell, col. 7, line 56 to col. 8, line 23)*; and

modifying the data structure *(i.e. modifies the data structure to reflect the responses, col. 8, lines 55-62)* to indicate that the second data is stored in the second image and storing the second data in the second image *(i.e. If it is for a write request, the routine continues to step 329 where it retrieves the data list for the address cell and removes all of the data cells from the list before the data cell corresponding to the write, col. 8, line 64 to col. 9, line 9)*.

Paulson does not specifically teach:

receiving a command to <u>create a backup of data stored in the data storage</u> <u>system</u>;

a second storage volume;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command; and

Flynn teaches:

receiving a command to access data to the data storage system *(i.e. In a virtual machine environment, a virtual machine ID associated with a virtual machine is sent to a storage controller along with a* <u>request to access a storage</u> <u>device</u>. *The virtual machine ID is used by the storage controller along with a path group ID to determine the scope of the access request. Thus, virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally, Summary)*;

a second storage volume *(i.e. Accordingly, a storage controller is operative to deny a storage device access request from a first virtual machine of a host processor responsive to a compare by the storage controller of a first virtual machine ID of the first virtual machine and a second virtual machine ID of a second virtual machine of the host processor. More specifically, the storage controller may be operative to receive a request for accessing a storage device; operative to compare a stored path group ID and a requesting path group ID associated with the request; operative to compare a stored virtual machine ID and a requesting virtual machine ID associated with the request; and operative to grant or deny the request responsive to the compares, col. 10, lines 24-37)*;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command *(i.e. the guest operating system performs device operations (read, write, etc.) on storage device 128, col. 8, lines 24-36)*.

It would have been obvious to one of ordinary skill of the art having the teaching of Paulson and Flynn at the time the invention was made to modify the system of Paulson to include the limitations as taught by Flynn. One of ordinary skill in the art would be motivated to make this combination in order to determine the scope of the access request in view of Flynn (Summary), as doing so would give the added benefit of virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally as taught by Flynn (Summary).

Paulson and Flynn do not explicitly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage system;

Tzelnic teaches this limitation *(i.e. <u>In response to the back-up request</u>, the method includes selecting a physical storage unit containing the logical data structure, and sending a backup command to the primary data storage subsystem for backing up the physical storage unit by copying a current version of the physical storage unit from primary storage in the primary data storage subsystem to the secondary data storage subsystem <u>to create a backup version</u> of the physical storage unit in secondary storage in the secondary data storage subsystem, col. 1, line 58 to col. 2, line 17)*.

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson, Flynn, and Tzelnic at the time the invention was made to

modify the request for accessing a storage system of Paulson, Flynn to include a

request for creating a backup of data in a storage system as taught by Tzelnic.

One of ordinary skill in the art would be motivated to make this combination in

order to restore the logical data structure in view of Tzelnic (col. 1, line 58 to col.

2, line 17), as doing so would give the added benefit of restoring the data rapidly

from backup after the storage system failure as taught by Tzelnic (col. 1, lines

33-55).

     As per claim 43, Paulson teaches a data storage system, comprising:

     a processing system *(See Fig. 1)*; and

     a memory coupled to the processing system, the memory storing

instructions, which when executed by the processing system *(See Fig. 1)*, cause

the processing system to perform the operations of:

     receiving a command to write data to the data storage system *(i.e. a*

second write request ("2w") to write the value 2, a second read request ("2f"), col.

4, lines 45-50);

     executing *(i.e. a first write request ("1w") to write the value 1, a first read*

("1r") request, col. 4, lines 45-50), for a first data, a fist input/output (I/O) process

(i.e. a first write request ("1w") to write the value 1, a first read ("1r") request, col.

4, lines 45-50) directed to a first storage volume, wherein the first storage volume

is not mirrored and the first I/O process begins at a first time which is prior to

receiving the command *(i.e. a second write request ("2w") to write the value 2, a*

second read request ("2f"), col. 4, lines 45-50);

creating a data structure, in response to the command, for at least a

second image, the second image storing changes to the first storage volume

occurring after receipt of the command *(i.e. it creates a new data cell, stores the*

data from the write request in the new cell, and appends the new data cell to the

end of the data list for this address cell, col. 7, line 56 to col. 8, line 23);

writing a second data *(i.e. When the second write request is received,*

another data cell would be created that held the value of 2, col. 7, line 56 to col.

8, line 23) directed to the second image as part of a second I/O process *(i.e. a*

second write request ("2w") to write the value 2, a second read request ("2f"), col.

4, lines 45-50) which begins after receipt of the command *(i.e. it creates a new*

data cell, stores the data from the write request in the new cell, and appends the

new data cell to the end of the data list for this address cell, col. 7, line 56 to col.

8, line 23); and

modifying the data structure *(i.e. modifies the data structure to reflect the*

responses, col. 8, lines 55-62) to indicate that the second data is stored in the

second image and storing the second data in the second image *(i.e. If it is for a*

write request, the routine continues to step 329 where it retrieves the data list for

the address cell and removes all of the data cells from the list before the data cell

corresponding to the write, col. 8, line 64 to col. 9, line 9).

Paulson does not fairly teach:

receiving a command to <u>create a backup of data stored in the data storage system</u>;

a second storage volume;

writing a second data directed to the second storage volume as part of a second I/O process which begins after receipt of the command; and

Flynn teaches:

receiving a command to access data to the data storage system *(i.e. In a virtual machine environment, a virtual machine ID associated with a virtual machine is sent to a storage controller along with a <u>request to access a storage device</u>. The virtual machine ID is used by the storage controller along with a path group ID to determine the scope of the access request. Thus, virtual machines of a single host processor may share the storage device while both preserving data integrity and performing optimally, Summary)*;

a second storage volume *(i.e. Accordingly, a storage controller is operative to deny a storage device access request from a first virtual machine of a host processor responsive to a compare by the storage controller of a first virtual machine ID of the first virtual machine and a second virtual machine ID of a second virtual machine of the host processor. More specifically, the storage controller may be operative to receive a request for accessing a storage device; operative to compare a stored path group ID and a requesting path group ID associated with the request; operative to compare a stored virtual machine ID and a requesting virtual machine ID associated with the request; and operative to grant or deny the request responsive to the compares, col. 10, lines 24-37)*;

writing a second data directed to the second storage volume as part of a

second I/O process which begins after receipt of the command *(i.e. the guest*

operating system performs device operations (read, write, etc.) on storage device

128, col. 8, lines 24-36).

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson and Flynn at the time the invention was made to modify the

system of Paulson to include the limitations as taught by Flynn. One of ordinary

skill in the art would be motivated to make this combination in order to determine

the scope of the access request in view of Flynn (Summary), as doing so would

give the added benefit of virtual machines of a single host processor may share

the storage device while both preserving data integrity and performing optimally

as taught by Flynn (Summary).

Paulson and Flynn do not explicitly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage

system;

Tzelnic teaches this limitation *(i.e. <u>In response to the back-up request</u>, the*

method includes selecting a physical storage unit containing the logical data

structure, and sending a backup command to the primary data storage

subsystem for backing up the physical storage unit by copying a current version

of the physical storage unit from primary storage in the primary data storage

subsystem to the secondary data storage subsystem <u>to create a backup version</u>

of the physical storage unit in secondary storage in the secondary data storage

subsystem, col. 1, line 58 to col. 2, line 17).

It would have been obvious to one of ordinary skill of the art having the
teaching of Paulson, Flynn, and Tzelnic at the time the invention was made to
modify the request for accessing a storage system of Paulson, Flynn to include a
request for creating a backup of data in a storage system as taught by Tzelnic.
One of ordinary skill in the art would be motivated to make this combination in
order to restore the logical data structure in view of Tzelnic (col. 1, line 58 to col.
2, line 17), as doing so would give the added benefit of restoring the data rapidly
from backup after the storage system failure as taught by Tzelnic (col. 1, lines
33-55).

As to claims 2, 16, 30, Flynn teaches the first storage volume is a first
virtual logical unit (VLUN) and the second storage volume is a second VLUN *(i.e.
a first virtual machine of a host processor responsive to a compare by the
storage controller of a first virtual machine ID of the first virtual machine and a
second virtual machine ID of a second virtual machine of the host processor, col.
10, lines 24-37).*

As to claims 6, 20, 34, Paulson teaches receiving a third data being
written to a data block of the second storage volume *(i.e. a third write request
("3w") to write the value 3, col. 4, lines 45-50);*
updating the data structure to indicate the data block is stored on the
second storage image *(i.e. modifies the data structure to reflect the responses,
col. 8, lines 55-62);* and

writing the third data to the data block on the second image *(i.e. FIG. 4E corresponds to the receipt of write request W3, indicating a request to write value W3 to memory address ADR1. FIG. 4E includes address cell 410 and data cells 420 and 430, col. 11, lines 5-28).*

As to claims 7, 21, 35, Paulson teaches updating comprises: determining whether the data block is stored on the first storage image *(i.e. FIG. 2 is a flow diagram for an implementation of the Request Tracker routine 205. The Request Tracker routine monitors read and write requests, and maintains a request data structure which contains the possible data which could be returned by the satisfaction of an accurate read request ... For each read request, the routine verifies that the read address is accurate, and if so it updates the corresponding address cell and its data list to reflect a pending read request. For each write request, the routine creates a new address cell if one does not yet exist for the address, and stores the data to be written in either a new or existing data cell for the address, col. 6, line 62 to col. 7, 18);* and

updating the data structure to indicate the data block is stored on the second storage image, if the data block is stored on the first image *(i.e. FIG. 2 is a flow diagram for an implementation of the Request Tracker routine 205. The Request Tracker routine monitors read and write requests, and maintains a request data structure which contains the possible data which could be returned by the satisfaction of an accurate read request ... For each read request, the routine verifies that the read address is accurate, and if so it updates the corresponding address cell and its data list to reflect a pending read request. For*

each write request, the routine creates a new address cell if one does not yet

exist for the address, and stores the data to be written in either a new or existing

data cell for the address, col. 6, line 62 to col. 7, 18).

As to claims 8, 22, 36, Paulson teaches examining a lookup table to

determine whether there is an entry associated with the data block, the lookup

table being associated with the second storage image *(i.e. FIG. 2 is a flow*

diagram for an implementation of the Request Tracker routine 205. The Request

Tracker routine monitors read and write requests, and maintains a request data

structure which contains the possible data which could be returned by the

satisfaction of an accurate read request … For each read request, the routine

verifies that the read address is accurate, and if so it updates the corresponding

address cell and its data list to reflect a pending read request. For each write

request, the routine creates a new address cell if one does not yet exist for the

address, and stores the data to be written in either a new or existing data cell for

the address, col. 6, line 62 to col. 7, 18); and

creating the entry associated with the data block if the entry does not exist

(i.e. FIG. 2 is a flow diagram for an implementation of the Request Tracker

routine 205. The Request Tracker routine monitors read and write requests, and

maintains a request data structure which contains the possible data which could

be returned by the satisfaction of an accurate read request … For each read

request, the routine verifies that the read address is accurate, and if so it updates

the corresponding address cell and its data list to reflect a pending read request.

For each write request, the routine creates a new address cell if one does not yet

exist for the address, and stores the data to be written in either a new or existing

data cell for the address, col. 6, line 62 to col. 7, 18).

As to claims 11, 25, 39, Paulson teaches receiving a request to read from

a data block on the second storage volume *(i.e. FIG. 2 is a flow diagram for an*

implementation of the Request Tracker routine 205. The Request Tracker routine

monitors read and write requests, and maintains a request data structure which

contains the possible data which could be returned by the satisfaction of an

accurate read request ... For each read request, the routine verifies that the read

address is accurate, and if so it updates the corresponding address cell and its

data list to reflect a pending read request. For each write request, the routine

creates a new address cell if one does not yet exist for the address, and stores

the data to be written in either a new or existing data cell for the address, col. 6,

line 62 to col. 7, 18);

determining whether the data block is stored in the first image or the

second image, based the data structure associated with the second storage

image *(i.e. FIG. 2 is a flow diagram for an implementation of the Request Tracker*

routine 205. The Request Tracker routine monitors read and write requests, and

maintains a request data structure which contains the possible data which could

be returned by the satisfaction of an accurate read request ... For each read

request, the routine verifies that the read address is accurate, and if so it updates

the corresponding address cell and its data list to reflect a pending read request.

For each write request, the routine creates a new address cell if one does not yet

exist for the address, and stores the data to be written in either a new or existing

data cell for the address, col. 6, line 62 to col. 7, 18);

reading the data block from the first image if the data block is stored in the

first image *(i.e. FIG. 2 is a flow diagram for an implementation of the Request*

Tracker routine 205. The Request Tracker routine monitors read and write

requests, and maintains a request data structure which contains the possible

data which could be returned by the satisfaction of an accurate read request …

For each read request, the routine verifies that the read address is accurate, and

if so it updates the corresponding address cell and its data list to reflect a

pending read request. For each write request, the routine creates a new address

cell if one does not yet exist for the address, and stores the data to be written in

either a new or existing data cell for the address, col. 6, line 62 to col. 7, 18); and

reading the data block from the second image if the data block is stored in

the second image *(i.e. FIG. 2 is a flow diagram for an implementation of the*

Request Tracker routine 205. The Request Tracker routine monitors read and

write requests, and maintains a request data structure which contains the

possible data which could be returned by the satisfaction of an accurate read

request … For each read request, the routine verifies that the read address is

accurate, and if so it updates the corresponding address cell and its data list to

reflect a pending read request. For each write request, the routine creates a new

address cell if one does not yet exist for the address, and stores the data to be

written in either a new or existing data cell for the address, col. 6, line 62 to col.

7, 18).

As to claims 12, 26, 40, Paulson teaches examining a lookup table to

determine whether there is an entry associated with the data block, the lookup

table being associated with the second storage image *(i.e. FIG. 2 is a flow*

diagram for an implementation of the Request Tracker routine 205. The Request

Tracker routine monitors read and write requests, and maintains a request data

structure which contains the possible data which could be returned by the

satisfaction of an accurate read request ... For each read request, the routine

verifies that the read address is accurate, and if so it updates the corresponding

address cell and its data list to reflect a pending read request. For each write

request, the routine creates a new address cell if one does not yet exist for the

address, and stores the data to be written in either a new or existing data cell for

the address, col. 6, line 62 to col. 7, 18).

As to claims 44, 45, 46, Flynn teaches the second I/O process is capable

of accessing the same data, via the second storage volume, as the first I/O

process *(i.e. a first virtual machine of a host processor responsive to a compare*

by the storage controller of a first virtual machine ID of the first virtual machine

and a second virtual machine ID of a second virtual machine of the host

processor, col. 10, lines 24-37).

Claims 4, 5, 9, 10, 13, 14, 18, 19, 23, 24, 27, 28, 32, 33, 37, 38, 41, 42 are rejected under 35 U.S.C. 103(a) as being unpatentable over Paulson et al. (US Patent No. 6,112,319), in view of Flynn et al. (US Patent No. 6,453,392), and Tzelnic et al. (US Patent No. 6,366,987), and further in view of Kakivaya et al. (US Patent No. 6,546,443).

As to claims 4, 18, 32, Paulson, Flynn and Tzelnic do not explicitly teach acquiring a lock from a lock mechanism before modifying the data structure to indicate that the second data is stored in the second image; and releasing the lock after storing the second data in the first image.

Kakivaya teaches:

acquiring a lock from a lock mechanism before modifying the data structure to indicate that the second data is stored in the second image *(i.e. a method and system for providing reader/writer synchronization services using interlocked operations, col. 3, lines 10-16)*; and

releasing the lock after storing the second data in the first image *(i.e. ReleaseWriterLock( ), col. 9, lines 41-52)*.

It would have been obvious to one of ordinary skill of the art having the teaching of Paulson, Flynn and Tzelnic, Kakivaya at the time the invention was made to modify the system of Paulson, Flynn and Tzelnic to include the limitations as taught by Kakivaya. One of ordinary skill in the art would be motivated to make this combination in order to provide reader/writer synchronization services using interlocked operations in view of Kakivaya (col. 3,

lines 10-16), as doing so would give the added benefit of providing a mechanism

for avoiding the deadlock problem as taught by Kakivaya (col. 3, lines 6-9).

   As to claims 5, 19, 33, Kakivaya teaches the lock mechanism is

maintained independent to the first and the second storage images *(i.e. a method*

and system for providing reader/writer synchronization services using interlocked

operations, col. 3, lines 10-16).

   As to claims 9, 23, 37, Paulson, Flynn and Tzelnic do not teach acquiring

a lock from a lock mechanism before the updating; and releasing the lock after

the writing.

   Kakivaya teaches:

   acquiring a lock from a lock mechanism before the updating *(i.e. a method*

and system for providing reader/writer synchronization services using interlocked

operations, col. 3, lines 10-16); and

   releasing the lock after the writing *(i.e. ReleaseWriterLock( ), col. 9, lines*

41-52).

   It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson, Flynn and Tzelnic, Kakivaya at the time the invention was

made to modify the system of Paulson, Flynn and Tzelnic to include the

limitations as taught by Kakivaya. One of ordinary skill in the art would be

motivated to make this combination in order to provide reader/writer

synchronization services using interlocked operations in view of Kakivaya (col. 3,

lines 10-16), as doing so would give the added benefit of providing a mechanism

for avoiding the deadlock problem as taught by Kakivaya (col. 3, lines 6-9).

As to claims 10, 24, 38, Kakivaya teaches the lock mechanism is

maintained independent to the first and the second storage images *(i.e. a method*

and system for providing reader/writer synchronization services using interlocked

operations, col. 3, lines 10-16).

As to claims 13, 27, 41, Paulson, Flynn and Tzelnic do not teach

acquiring a lock from a lock mechanism before the determining; and releasing

the lock after the reading.

Kakivaya teaches:

acquiring a lock from a lock mechanism before the determining *(i.e. a*

method and system for providing reader/writer synchronization services using

interlocked operations, col. 3, lines 10-16); and

releasing the lock after the reading *(i.e. ReleaseWriterLock( ), col. 9, lines*

41-52).

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson, Flynn and Tzelnic, Kakivaya at the time the invention was

made to modify the system of Paulson, Flynn and Tzelnic to include the

limitations as taught by Kakivaya. One of ordinary skill in the art would be

motivated to make this combination in order to provide reader/writer

synchronization services using interlocked operations in view of Kakivaya (col. 3,

lines 10-16), as doing so would give the added benefit of providing a mechanism for avoiding the deadlock problem as taught by Kakivaya (col. 3, lines 6-9).

As to claims 14, 28, 42, Kakivaya teaches the lock mechanism is maintained independent to the first and the second storage images *(i.e. a method and system for providing reader/writer synchronization services using interlocked operations, col. 3, lines 10-16).*

Response to Arguments

1.   Applicant's arguments with respect to claims 1-2, 4-16, 18-30, 32-46 have been considered but are moot in view of the new ground(s) of rejection.

Furthermore, in response to Applicant's arguments with respect to Paulson fails to describe receiving a command to create a backup of data stored in a data storage system as required by the amended claims, the Examiner respectfully notes that the term "a command to create a backup" is not supported by the specification (See [0004], Specification), the specification merely states a command to preserve data. Thus, the term "preserve" could be equivalent to a write operation data to a disk.

Accordingly, Paulson teaches receiving a command to write data to the data storage system *(i.e. a second write request ("2w") to write the value 2, a second read request ("2f"), col. 4, lines 45-50).*

Flynn teaches receiving a command to access data to the data storage system *(i.e. In a virtual machine environment, a virtual machine ID associated*

with a virtual machine is sent to a storage controller along with a <u>request to</u>

<u>access a storage device</u>. The virtual machine ID is used by the storage controller

along with a path group ID to determine the scope of the access request. Thus,

virtual machines of a single host processor may share the storage device while

both preserving data integrity and performing optimally, Summary).

Paulson and Flynn do not explicitly teach:

receiving a command to create <u>a backup of data stored in the</u> data storage

system;

Tzelnic teaches this limitation *(i.e. <u>In response to the back-up request</u>, the*

method includes selecting a physical storage unit containing the logical data

structure, and sending a backup command to the primary data storage

subsystem for backing up the physical storage unit by copying a current version

of the physical storage unit from primary storage in the primary data storage

subsystem to the secondary data storage subsystem <u>to create a backup version</u>

of the physical storage unit in secondary storage in the secondary data storage

subsystem, col. 1, line 58 to col. 2, line 17).

It would have been obvious to one of ordinary skill of the art having the

teaching of Paulson, Flynn, and Tzelnic at the time the invention was made to

modify the request for accessing a storage system of Paulson, Flynn to include a

request for creating a backup of data in a storage system as taught by Tzelnic.

One of ordinary skill in the art would be motivated to make this combination in

order to restore the logical data structure in view of Tzelnic (col. 1, line 58 to col.

2, line 17), as doing so would give the added benefit of the data should be

restored rapidly from backup after the storage system failure as taught by Tzelnic

(col. 1, lines 33-55).

Conclusion

Applicant's amendment necessitated the new ground(s) of rejection

presented in this Office action. Accordingly, **THIS ACTION IS MADE FINAL**.

See MPEP § 706.07(a). Applicant is reminded of the extension of time policy as

set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire

THREE MONTHS from the mailing date of this action. In the event a first reply is

filed within TWO MONTHS of the mailing date of this final action and the advisory

action is not mailed until after the end of the THREE-MONTH shortened statutory

period, then the shortened statutory period will expire on the date the advisory

action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be

calculated from the mailing date of the advisory action. In no event, however, will

the statutory period for reply expire later than SIX MONTHS from the date of this

final action.

Any inquiry concerning this communication or earlier communications from the
examiner should be directed to Miranda Le whose telephone number is (571)
272-4112. The examiner can normally be reached on Monday through Friday
from 10:00 AM to 6:00 PM.
        If attempts to reach the examiner by telephone are unsuccessful, the
examiner's supervisor, James K. Trujillo, can be reached on (571) 272-3677.
The fax number to this Art Unit is (571)-273-8300.

        Any inquiry of a general nature or relating to the status of this application
should be directed to the Group receptionist whose telephone number is (571)
272-2100.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see http://pair-direct.uspto.gov. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

/Miranda Le/

Primary Examiner, Art Unit 2169